

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Antonio LAIN et al.

Title: MANAGEMENT OF SECURITY KEY DISTRIBUTION

Appl. No.: Unassigned

Filing Date: 10/29/2003

Examiner: Unassigned

Art Unit: Unassigned

CLAIM FOR CONVENTION PRIORITY

Commissioner for Patents
PO Box 1450
Alexandria, Virginia 22313-1450

Sir:

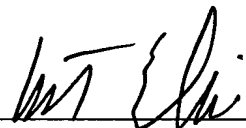
The benefit of the filing date of the following prior foreign application filed in the following foreign country is hereby requested, and the right of priority provided in 35 U.S.C. § 119 is hereby claimed.

In support of this claim, filed herewith is a certified copy of said original foreign application:

- Great Britain Patent Application No. 0225407.6 filed 10/31/2002.

Respectfully submitted,

Date October 29, 2003

By  _____

FOLEY & LARDNER
Customer Number: 22879
Telephone: (202) 672-5485
Facsimile: (202) 672-5399

William T. Ellis
Attorney for Applicant
Registration No. 26,874



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.


In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated 29 January 2003



Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference	200205659-1 GB		
2. Patent application number <i>(The Patent Office will fill in this part)</i>	0225407.6		31 OCT 2002
3. Full name, address and postcode of the or of each applicant <i>(underline all surnames)</i>	Hewlett-Packard Company 3000 Hanover Street Palo Alto CA 94304, USA Delaware, USA		
Patents ADP number <i>(if you know it)</i> If the applicant is a corporate body, give the country/state of its incorporation	496588001 7563083001		
4. Title of the invention	Management of security key distribution		
5. Name of your agent <i>(if you have one)</i>	Bruce G R Jones Hewlett-Packard Ltd, IP Section Filton Road, Stoke Gifford Bristol BS34 8QZ		
"Address for service" in the United Kingdom to which all correspondence should be sent <i>(including the postcode)</i> Patents ADP number <i>(if you know it)</i>			
6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and <i>(if you know it)</i> the or each application number	Country	Priority application number <i>(if you know it)</i>	Date of filing <i>(day / month / year)</i>
7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	Number of earlier application		Date of filing <i>(day / month / year)</i>
8. Is a statement of inventorship and of right to grant of a patent required in support of this request? <i>(Answer 'Yes' if:</i>	Yes		
a) any applicant named in part 3 is not an inventor, or b) there is an inventor who is not named as an applicant, or c) any named applicant is a corporate body. See note (d))			

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description

18

Claim(s)

2

Abstract

1

Drawing(s)

3 +3

8

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

1

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

1

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

Fee Sheet

—

11. I/We request the grant of a patent on the basis of this application.

Signature

Bruce G R Jones

Date 31/10/2002

12. Name and daytime telephone number of person to contact in the United Kingdom

Tony Judd

Tel: 0117-312-8026

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

MANAGEMENT OF SECURITY KEY DISTRIBUTION

The present invention relates to the management of security key distribution, most typically within a community of anonymous users all of whom are sharing a secure service, which may be the provision of content or the provision of a resource for example.

One example of such a situation is one in which a large number of anonymous users subscribe to a service providing shared content which is updated on a regular basis.

To protect the interests of the provider of the content such content is usually distributed to bona fide users (also known herein as "subscribers") in an encrypted form. This prevents non-subscribers from gaining access to the content and thereby diminishing its financial value to existing subscribers, and thus ultimately the pecuniary advantage that may be obtained by the provider. In one example each subscriber is given a key which they may use to decrypt content; to protect the interests of subscribers, such a key should ideally neither identify them nor enable such identification. However, it has long been established that managing the provision and maintenance of security keys to a large group of anonymous subscribers is difficult. For example, one way in which both the anonymity of the subscribers may be preserved and the provision of keys may be made simple is to give each user the same security key, however this has negative implications on the security offered by such a single key. In an alternative key management method, each user is issued with a key which is unique, at least within the provider of the content, but which does not identify the subscriber, and which functions to decrypt only content sent to that subscriber. In such a scenario, upon lapsing of the subscription, it is possible to invalidate this unique key by ceasing to make content available in a form which is decryptable using the key issued to the now-lapsed subscriber. However owing to the manner in which such keys are generated in the vast majority of instances, the revocation of even such unique keys from lapsed subscribers requires a reconfiguration of all other subscriber's keys, at least to some extent, and eventually, when sufficient keys have been invalidated, the need to reissue keys in their entirety to all subscribers.

A first aspect of the present invention relates to the revocation of unique keys from lapsed subscribers, and more particularly the basis upon which it is decided to revoke such keys. According to a first aspect of the present invention there is provided a method of managing security keys provided to users of a service comprising the steps of:

- issuing a security key to a first user eligible to receive the service;
- monitoring the first user's status to establish whether the first user is eligible to receive the service;
- establishing, in accordance with a policy, a first value associated with invalidation of the first user's key, and a second value associated with providing the service to an ineligible user, and if the second value exceeds the first value, invalidating the key.

In a preferred embodiment the policy is based on economic grounds, so that instead of invalidating a key simply on a contractual basis because a subscription has lapsed for example, the cost to the provider of doing so is assessed, and invalidation takes place at an optimised instant in time from the point of view of the provider. Accordingly the first value preferably represents the cost to the provider of invalidating a key, and the second value represents the cost of providing the service to an ineligible user. The first value may typically include what may be thought of as consequential costs, including one or more of: the cost to the provider of disrupting the provision of the service as a result of having to reconfigure all the other issued keys to some extent, and the likelihood that the invalidation of a key will trigger the need to reissue all keys in their entirety.

The second value preferably takes into account aggregated costs of providing network capacity and server capacity to all currently ineligible users, and the economic effects of dilution of value of the service to remaining users, such as for example any consequentially increased tendency to pay subscriptions late, for example.

Frequently different levels of service are offered, and under different commercial terms, such as length of a subscription paid for in advance (and privileges associated with that for example). A second and independent aspect of the present invention relates to an appreciation of the fact that, where security keys are generated in a

structured manner, such as a hierarchy for example, it may be advantageous to take into account user characteristics, and to allocate keys from the hierarchy on the basis of such characteristic.

5 According to a second aspect of the present invention, there is provided a method of managing security keys generated from an ancestral hierarchy and used to provide selective access to provision of a service, wherein invalidation of a key necessitates reconfiguration of each other key within the hierarchy to the extent another key and an invalidated key share common ancestry, the method comprising the steps of:

10 defining, upon the basis of a predetermined policy, at least two groups of users of the service to whom keys have been issued;

allocating within the hierarchy a distinct domain for each group of users; and

issuing keys to users from domains within the hierarchy upon the basis of their grouping.

15

According to one embodiment, a group of users who have contracted to a high level of service and are therefore perceived to be valuable to the provider are allocated keys from a first domain within the hierarchy, an important characteristic of which is that keys from the first domain share fewer ancestors with keys from other domains of the hierarchy than those other keys share with each other. Consequently, when a key is
20 invalidated from a domain other than the first domain, the keys of users from the first group require less reconfiguration than the keys from any other domain, so that the most valuable users are inconvenienced the least.

25 Embodiments of the present invention will now be described, by way of example, and with reference to the accompanying drawings, in which:

Fig. 1 is a schematic illustration of a network in which a secure resource is to be provided to subscribers;

30

Fig. 2 is a schematic illustration of an ancestral hierarchy of security keys for securing the provision of a service;

Fig. 3 is a schematic illustration of a process of requesting a key and then using the key to retrieve secure content; and

Fig. 4 is a schematic illustration of a further ancestral hierarchy of security keys, and the manner in which domains within which such a hierarchy may be defined to correspond to user groups.

Referring now to Fig. 1, a service is provided by a server computing entity 10 to a plurality of subscribers 20, who in the present example are in a client (computing entity) relationship with the server 10. It should be noted that the present invention is applicable without limitation to the nature of the relationship between the provider of the resource and the subscribers to its provision, and may thus for example find application in virtually any other architecture and/or relationship, such as peer to peer networks. The present illustrated architecture and relationship has been chosen merely for simplicity of illustration of the principles underlying the present invention. In the example of Fig. 1 the resource provided to subscribers is content 100 in documentary form, stored on the server 10. However, once again the present invention is equally applicable to the provision of any resource which is of value to subscribers.

To protect the economic value of the content the provider prohibits assimilation thereof by persons who are not in possession of a key, while issuing such a key to each subscriber. The prohibition may operate at any one (or more) of the stages of a process which includes the steps of: retrieving the content from the server via the network; saving the content on a subscriber's client machine; and consuming the content (e.g. in the case of visually assimilable content, reading it). The general examples given above are applicable to the provision of a resource in the form of content; in the case of the provision of some other resource, for example the use of a particular hardware element on the server machine, the prohibition is likely to operate in a somewhat different manner. Allocation and management of keys is performed at the server by a key management program KMP. The precise manner of operation of both prohibition upon the availability of a resource, and use of a key to provide exception to such prohibition are not germane to the present invention, and will not therefore be discussed in any detail. More complete information may however be

obtained from "Key Management for Multicast: Issues and Architectures; D. Wallner, E. Hardner & R. Agee, available online at:

<http://www.ietf.org/rfc/rfc2627.txt> the contents of which are hereby incorporated by reference.

5

Technical issues which must be taken into account in management of the provision of keys to subscribers result not merely from solely technical considerations; commercial considerations similarly give rise to technical issues. One such commercial consideration is the desire of many subscribers to remain anonymous. There are
 10 many possible reasons for this. For example, where the resource is content, the nature of the content may be such that there is a degree of moral opprobrium associated with persons consuming it, such as for example sexually explicit material. Alternatively, there may be commercial reasons for wishing to remain anonymous, for example where a subscriber to the content is a commercial undertaking, knowledge of the
 15 nature of the content to which they are subscribing may provide competitors with useful information regarding their operations or future intentions. It is therefore necessary in such a situation for the key which is distributed to each subscriber to be intrinsically incapable of identifying the subscriber. A further commercial consideration is the frequent commercial need to provide for differing levels of
 20 subscription, corresponding to different levels of content provision, or different levels of service associated with such content provision. For example "Gold" users of an online news service are able to gain access to web pages which are updated every hour, whereas "Bronze" users can only access pages which are updated once per day. Preferably therefore, the key structure should enable a provider who offers differing
 25 levels of service to reflect these service levels in any key management activities.

One method in which encryption keys may be generated is illustrated in highly abstracted and simplified form in Fig. 2. In this example, encryption keys are generated in an ancestral hierarchy, which here has the form of a binary tree, although
 30 neither the binary nature of the tree, nor the tree-like architecture is essential. Features of such a key structure which are relevant to an understanding of the present invention are:

1. Each key may be used to generate two further keys, so that ultimately all keys in the tree are related to an end or root key A0.
2. Due to the binary nature of the tree, the size (e.g. in terms of the numbers of characters) of the keys in a given generation d is twice that of the keys in its parent generation $d-1$.
3. Each key will intrinsically indicate its provenance within the tree in the form of a path down to that key from the root key.
4. Issue of a key in a given generation d to a subscriber compromises the security provided by any keys directly related to the issued key lower down the hierarchy, thus rendering all directly descended keys in generations $(d + n)$, where n is an integer, redundant.
5. Use of a key located at a generation d within the tree implicitly includes the use of all keys of higher generations within the tree on a path down from the root key.
6. Invalidation of any key requires reconfiguration of all other keys in the hierarchy to the extent that they share common ancestry.

A number of consequences flow from these characteristics. Firstly, and most obviously, the root A0 cannot ever be issued to a subscriber, since this would compromise the security provided by all of the other keys in the hierarchy. Secondly, and following on from this, the higher the generation of a given key, the more costly it is to issue that key to a subscriber in terms of the number of lower generation keys which are compromised as a result and are therefore redundant within the hierarchy - either for use as issued keys to subscribers or for generating descendent keys (since any descendent keys would likewise be compromised) . Thirdly, the length of the key will indicate to its provenance in terms of generation within the hierarchy.

In the tree of Fig. 2 a total of only four generations have been illustrated. In practice a tree of this type is likely to have many more than four generations in total, with the total number of generations being denoted by a variable k , giving, for a binary tree a total of $2^1 + 2^2 + \dots + 2^k$ keys, although not all of these can be allocated because of the characteristic referred to at (4) above. In the example of Fig. 2, a square around the identifier for a key is indicative of the fact that that key has been allocated to a subscriber. Thus the key D2 has been allocated to a subscriber, which has the

consequence that keys G3, H3 and M4 to P4 - a total of six keys in a tree with only four generations - are unuseable as a result. In fact, were a key of such an elevated generation be allocated in a practical size of tree the effect on the number of "lost" keys would be very large indeed, equal to approximately a quarter of the total keys (since the allocated key is in the second generation containing only four keys), and so it is unlikely that in practice such a key would be allocated. Because of the nature of the keys and their generation, the use of any key within the tree implicitly incorporates the use of all keys of higher generations on a path up to the root key. Thus, as the generations of the tree are descended, keys of lower generations will tend to share more common ancestor keys of higher generations, and this has implications both for the level of security offered by the use of such keys in that they are easier to break, and also for their invalidation, since this will have ramifications on the security provided by still-active keys sharing the same higher-generation ancestors. Thus allocation of a relatively high generation key does have advantages, in that its invalidation causes ramifications for fewer peer generation keys, and the key is short, meaning that it is easy to transmit and store (which takes place on the subscriber's machine). Conversely, a further disadvantage related to the use of keys from lower generations within the tree is that these keys become increasingly large and unwieldy with increasing generations.

The use of a key in the process of distributing secure content to a subscriber will now be described briefly with reference to Fig. 3. The process is initiated by transaction 300 with receipt by the subscriber of a request 302 for a subscription from a would-be subscriber. The request at this stage is likely to include some form of payment, or promise to pay. Given the need for anonymity of subscribers, the processing of any payment at the server is likely to be dealt with first, and successful payment processing will simply yield an indication to a key management program 110 that the requesting subscriber is entitled to a given level of content provision, e.g. "GOLD" for a given period of time. The key management program then allocates a key to the requesting subscriber at operation 304, and this key is sent back to the subscriber at transaction 306 as part of a cookie 308, which includes the key, here represented as #D2, and a time to live DDMMYYYY for this key, this being the date when the subscription paid for in the payment step expires. It should be noted at this point that, unlike the majority of cookies usually passed to clients from servers, in the present

example the cookie 308 does not give any indication of the identity of the requesting subscriber, and other than during the payment process, which is performed entirely separately to the key allocation process, no data identifying the subscriber has been received by the server. NB In a modification, payment may be made by the subscriber to a trusted third party, who, once payment has been made, then simply passes data indicating the manner in which the client may be contacted to the server for the allocation of a key to that subscriber. Upon receipt of the key #D2, the subscriber then stores this key securely at operation 310. When the subscriber then wants to gain access to the content 100 in respect of which they have subscribed, the cookie 308 containing the key #D2 is retrieved at operation 312 and sent to together with a request to access the content in transaction 314. When such a request is received, the server extracts the key #D2 from the cookie, and using the key #D2 authenticates the request at operation 316.

The process of authentication does not include any process in which the key submitted by the subscriber is mapped to an identity for them, since this would inherently compromise their anonymity. Rather, the process simply involves determining whether the key is a genuine key (i.e. one generated from the tree), and whether the level of content (for example, as mentioned above, hourly update rather than daily, for example) indicated within by the key corresponds to that being requested by the user. Thus, presentation of an authenticated key is *per se* verification of entitlement to the content, meaning that responsibility for secure retention of the key is entirely the subscriber's, since a key appropriated from the subscriber by an unscrupulous third party would enable that third party to gain access to the content. A further consequence is that the subscriber bears the entire burden of responsibility for the protection of their anonymity *vis-à-vis* their subscription to the content, in that by gaining access to their machine, it is possible to ascertain their identity, and a mapping of their identity to their key exists implicitly on their machine since this is where the key is stored.

The specific process of authentication involves firstly location of the key #D2 within the tree of Fig. 2. As mentioned above, each key includes an indication of the its address within the tree in the form of a path from the root key A0, so that in the present example, the key #D2 will implicitly contain data indicative of the path

A0-B1-D2. Once the node or address D2 within the tree has been located, the key #D2 is simply matched with the key at that address, and if they are the same then the key #D2 is assumed to be genuine. Authentication of the level of content provision may be performed in a number of ways. In one embodiment a mapping is made of each key issued from the tree to the level of content provision for which the subscriber has paid, so that upon receipt and authentication of the key, the authenticated key is then mapped to the content level, and the content level indicated by this mapping is matched with the requested content level. If both the key and the requested content level are authenticated at operation 316, the server then retrieves the requested content and encrypts it with the subscriber's key at operation 317, before dispatching the encrypted content to the subscriber at transaction 318. The subscriber then retrieves their key from secure storage, and uses the key to decrypt the content at operation 320. If, at some later stage, the subscriber no longer wishes to subscribe, they may request cancellation of their subscription at transaction 322, and following receipt of such a cancellation request, the server inactivates the key to prevent the now-cancelled subscriber from gaining access to the content without paying for it.

The above description of both the tree method of generating such keys, and the simplified scenario of the use of a key thus generated is both simplified and incomplete (as mentioned above a fuller explanation of this being provided in the document referenced by Wallner et al), and serves merely to provide sufficient information for an understanding of the context of the present invention, which, in imprecise terms, may be thought of as relating to the management of key allocation and maintenance in a commercial context.

A first aspect of the present invention relates to the management of keys in the event that a subscriber (or put in more general terms, an eligible or *bona fide* user) becomes ineligible, for example as a result of cancellation of a subscription (whether this is an active event, or by virtue of an existing subscriber failing to renew or pay for the next periodic subscription). Referring again to Fig. 2, when a subscriber to whom key K4 (this key being an allocated key as indicated by the square around it has been) has been allocated wishes to cancel their subscription, inactivation of their key has a number of consequences. For example, if the same subscriber subsequently wishes to

- renew their subscription shortly after cancelling it (for example in the event that they forgot to pay, or were late paying due to financial constraints), a new key will have to be issued to them. Given that the tree has only a finite number of keys, cancellation and reissue of a new key to the same subscriber will reduce by two the number of
- 5 keys available for use. The ultimate consequence of this is that eventually, all allocatable keys within the tree are used up, a new tree will have to be generated, and new keys reissued to existing subscribers. This is both expensive for the provider, and irksome for the subscribers.
- 10 Furthermore, and notwithstanding this, invalidation of any key results in a need for at least a degree of reconfiguration of any other key, to the extent it shares any of its ancestors, and then transmission of such reconfigured keys to the subscribers to which such keys have been allocated. More specifically, keys which only share root key A0 as an ancestor will require only one reconfiguration (and corresponding transmission
- 15 of such a reconfigured key), while those sharing two generations of ancestor require two reconfigurations and therefore two transmissions of reconfigured keys, and so on. It is thus readily apparent that invalidating any key has far reaching consequences. In the specific example of Key K4, its ancestral path is F3-C2-B1-A0, and so it shares:
- 20 - ancestral keys: F3, C2, B1 and A0 with key L4, which therefore must have all four reconfigured keys redistributed to it for its (i.e. L4's) complete reconfiguration;
- ancestral keys C2, B1 and A0 with keys I4, L4 and E3 each of which must therefore receive configured keys for keys C2, B1 and A0;
- ancestral keys B1 and A0 with keys D2, G3, H3, and M4 to P4, each of which must
- 25 therefore receive reconfigured keys for keys B1 and A0;
- root key with all the remaining keys, all of which must therefore receive reconfigured root keys.

In addition to this cost, and as referenced above, the allocation of any key to a

30 subscriber has an opportunity cost associated with it, in the form of the number of descendant keys which are redundant as result; the higher the generation of the allocated key, the greater the opportunity cost. Cancellation of a key may thus also be thought of as realising the associated opportunity cost: for example cancelling key #D2 amounts to cancellation of a total of seven keys (from node D2 and its six

at least part of all the other keys in the tree to the extent that they share ancestral keys with D2, and distribution of these reconfigured key elements their subscribers.

Balanced against this cost to the provider, is the economic loss they will suffer as a consequence of failing in fact to cancel the subscription by invalidating the key #D2. Elements contributing to this loss include the contribution to the total cost of providing the necessary network and server capacity, the economic effect of dilution of content value (including the potential exploitation of any leniency exhibited by the provider to late payment, for example), but zero marginal cost with regard to the creation or storage of the content, since both these operations must be conducted regardless. It is thus possible to create a policy in which all of these factors are taken into account, suitably weighted to reflect the specific circumstances (or indeed personal preferences) of the provider, so that on each occasion a quantitatively based decision may be made with regard to the invalidation of a key may be made.

One such policy provides for a continual decision making process on each occasion that a subscription lapses, and which takes into account the aggregated lapsed subscriptions at that instant in time to establish whether the costs a provider will incur as a result of invalidating a given subscriber's key are greater than costs to the provider of maintaining an unpaid-for service for all lapsed subscribers whose key are not invalidated at that instant in time; when the latter becomes greater than the former, all outstanding keys for lapsed subscriptions are invalidated. The manner in which this policy operates is shown in more detail below.

1. The cost to the provider of maintaining an unpaid-for service to lapsed subscribers:

Economic Dilution Cost (E_c) + Cost of Network Capacity (N_c)

The Network Capacity Cost is a widely varying cost which depends greatly upon the nature of the service provided. Thus for example in the case of content for which a high service level subscription has previously been paid, the content may include video streaming and other high data-rate transmission content items, in which case the cost of providing the content service is likely to be relatively high, whereas low level

content service is relatively inexpensive to provide. As with other costs, the true cost to the provider is the aggregated cost due to all instant lapsed subscribers.

The network capacity cost can be quantified as follows:

5

$$\sum_{i=1,j} (n/w \text{ usage}_i) \times (\text{ServerRate}_i + \text{NetworkRate}_i)$$

where: j is the contemporaneous number of lapsed subscribers;

10

$n/w \text{ usage}$ is the average number of bits per second a user i consumes per unit time; and

ServerRate and NetworkRate are the charging rates for server and network capacity (in terms of bits per second processes or transmitted respectively) per unit time.

15 The Economic Dilution Cost is quantified as follows:

$$\sum_{i=1,j} [X_{t_i} \times R \times P_i (\text{ForceResub})] - [P_i (X_t) \times P_i (\text{VolResub}) \times R]$$

20

where: X_{t_i} is the extra time which a subscriber i will enjoy for free as a result of failing to invalidate their key;

R is the subscription Rate per unit time;

$P_i (\text{ForceResub})$ is the probability that a subscriber i will re-subscribe as soon as they are forced to do so by invalidation of their key;

25

$P_i(X_t)$ is the probability of the subscriber receiving extra time – i.e. whether the tree will be rebuilt before the subscriber gets extra time;

$P_i(\text{VolResub})$ is the probability that subscriber i will re-subscribe voluntarily as a result of being given extra time.

Thus the expression:

30

$$[X_{t_i} \times R \times P_i (\text{ForceResub})]$$

represents the amount of money which is lost due to failing to force a lapsed subscriber to re-subscribe immediately, and the expression:

$$[P_i (X_t) \times P_i (\text{VolResub}) \times R]$$

5

represents the amount of money which is gained by a lapsed subscriber re-subscribing as a result of the liberal attitude to their lapsed subscription.

2. Cost of invalidating Lapsed Customer's Key

10

This is equal to:

Consequential Cost of breach of existing Service Level Agreements (SLA_{cost}) as a result of having to reissue keys + Cost of loss of consequentially disaffected customers (POC_{cost}) as result of key reissue + Key generation costs

15

$$SLA_{\text{cost}} = P(\text{OoS}) \times \text{NoSubscribers} \times \text{PenaltyCost}$$

Where: $P(\text{OoS})$ is the probability of key redistribution causing the service to drop below agreed levels (e.g. in the case of news service, updates less frequently than have been agreed);
 NoSubscribers is the total number of current subscribers;
 PenaltyCost is the cost for each subscriber of the service dropping below the agreed levels.

25

$$POC_{\text{cost}} = P(\text{depart}) \times \text{NoSubscribers} \times \text{SubscriberLifetimeValue}$$

where: $P(\text{depart})$ is the probability a customer will cancel their subscription as a result of the inconvenience of key regeneration;
 NoSubscribers is the total Number of subscribers
 $\text{SubscriberLifetimeValue}$ is the estimated future revenue each subscriber will yield to the provider

30

NB in the case of each of SLA_{cost} and POC_{cost} the parameters given above are for average values. More accurate calculations can take individual values into account by summation if desired.

$$5 \quad KEY_{costs} = [KeyR/genCost_i \times DistCost_i] + [P(T/TreeR/gen) \times DistCost_{T/Tree}]$$

where: $KeyR/genCost_i$ is the cost of regenerating a key for a subscriber i ;
 $DistCost_i$ is the cost of distribution necessary when regenerating a key for subscriber i ;

10 $P(T/TreeR/gen)$ is the probability of the invalidation of a key causing the need to regenerate the entire tree;

$DistCost_{T/Tree}$ is the cost of distributing keys to all subscribers upon regeneration of the tree.

15 Where probabilities have been used in these calculations, typically they are probabilities which are obtained empirically from historical data.

As referenced briefly above, one important aspect of this policy is that on each occasion a subscription lapses, the policy is applied to determine for that individual
 20 subscription, whether it is economically advantageous to invalidate the key of that subscriber. In this connection it should be noted that the potential costs of failing to invalidate the key of that subscriber are calculated by aggregating the cost of maintaining unpaid-for service to all lapsed subscribers whose keys have not been
 25 individual key under consideration. In the event that the individual key is invalidated, the still-current keys of other lapsed subscribers are not however invalidated. For this reason a preferred feature of this policy is the periodic re-calculation in accordance with the policy of for each still-current key, since the same calculation performed
 30 subsequently may well yield a different result in view of subsequent events. For example, in a situation where the key of the first of ten lapsed subscribers was not invalidated when the policy was first applied, this may have been because the aggregated contribution to the cost of maintaining unpaid-for service to the other nine lapsed subscribers was not taken into account (because temporally their subscription

lapsed subsequently, and so these costs were not present at that time); a subsequent recalculation may therefore yield a different result.

Eventually, notwithstanding the ameliorative effects of the policy elucidated above, the tree will have to be rebuilt in its entirety. A further independent aspect of the present invention relates to the manner in which keys are allocated from such a tree, and provides that this is performed on the basis of commercial subscriber characteristics. Thus for example, in accordance with one model of policy, a subscriber who has paid in advance for a whole year's subscription is economically valuable to the provider, and so a subscriber who it is desirable to keep happy. It may therefore be economically advantageous to allocate such a customer, and customers of a similar commercial value, keys of a relatively high generation, and from a part of the tree which is not affected by the invalidation of lower generation key and the consequential reconfiguration which must take place beyond the reconfiguration of the root key. Less economically valuable users will be allocated keys from generations and parts of the tree which are increasingly susceptible to increasing reconfiguration as a result of the invalidation of other keys in the tree.

Referring to Fig. 4, a tree similar to that of Fig. 2 is created for existing subscribers following expiry of the useful life of a previous tree. The provider is at this point able to identify amongst his existing subscribers at least the following groups: Gold – these are subscribers to a high level of service, who have a high economic value to the provider, and who have for example paid large subscription fees in advance; Silver – subscribers of a lesser economic value than Gold who are more inclined to cancel their subscription, but nonetheless subscribe to a significant service level; Bronze – largely low service level subscribers who retain valid subscriptions for only a short time period; and Churn – these are subscribers who subscribe for trial periods and/or to the lowest service level. To avoid aggravating the Gold subscribers with unnecessary reconfigurations of their key, these subscribers are allocated keys from a domain provided by one half of the tree structure, which is reserved in its entirety for them. Silver users are allocated keys from a domain provided by a quarter of the illustrated tree structure and in the half of the tree structure not allocated exclusively to Gold subscribers, and this quarter of the tree is likewise allocated exclusively to them. Bronze subscribers are allocated an domain provided by an eighth of the tree

structure one generation lower than the silver subscribers, with Churn subscribers being allocated a domain provided by one sixteenth of the tree one generation lower than Bronze users. The consequences of this commercially based architecture in terms of key regeneration are as follows. In the event that a subscription from the

5 Churn domain subscription lapses, all other Churn subscribers will have to have at least keys G3, D2, B1 and A0 reconfigured. Keys in the Bronze domain will require reconfiguration of keys D2, B1 and A0; keys in the Silver domain will require reconfiguration of keys B1 and A0, while keys in the Gold domain will require configuration of key A0. It follows that the consequences of frequent lapses in

10 subscription amongst the most fluid group of subscribers in terms of key regeneration will fall most heavily upon them, since other subscribers of this group will share most ancestral keys, and will therefore need the largest number of key reconfigurations. By contrast, because Gold subscriber's keys are from a domain another half of the tree entirely, which is exclusively reserved for them, any lapsed subscription which occurs

15 in the other half of the tree will only ever require Gold subscribers to reconfigure the root key A0.

In an alternative application of this architecture, a different commercial perspective may yield an opposite outcome. For example, where a group of subscribers have paid

20 a long term subscription, or are contractually bound to such a subscription, it may be desirable to allocate keys to them from a domain where they will have to undergo frequent key reconfigurations when another subscriber's key is invalidated. This will be inconvenient, but because they are already committed to a long term subscription, will not result in any short to medium term loss of revenue by the provider. Keys

25 from a domain in which subscribers do not suffer the inconvenience of large-scale key reconfigurations (e.g. the domain allocated to Gold subscribers in Fig. 4) may then be reserved for short term subscribers who are more mobile, in an attempt to retain their subscription.

30 In a further aspect of the present invention a decision can be made by the provider on the basis of policy whether to secure content to a subscriber at all. For example, in a situation where a subscriber has committed only to a trial subscription period, and to only a low level of content with a relatively short commercial life (e.g. in the case of a news service, where contemporaneous content is the *sine qua non* of the service), it

may be the case that the risk and potential consequences of misappropriation of the content provided to this user by third parties are such that it is not commercially worth cost, in terms of key allocation and management, of securing the content at all, and so no key is issued. In such a situation the user will most preferably be unaware of this

5 and a placebo key issued to the user will simply identify to the provider that no authentication is required prior to provider the level of content specified therein.

CLAIMS

1. A method of managing security keys generated from an ancestral hierarchy and used to provide selective access to provision of a service, wherein invalidation of a key necessitates reconfiguration of each other key within the hierarchy to the extent another key and an invalidated key share common ancestry, the method comprising the steps of:
 - defining, upon the basis of a predetermined policy, at least two groups of users of the service to whom keys have been issued;
 - allocating within the hierarchy a distinct domain for each group of users; and
 - issuing keys to users from domains within the hierarchy upon the basis of their grouping.
2. A method according to claim 1 wherein the policy provides that users are grouped according to their perceived value to a provider of the service.
3. A method according to claim 2 wherein a first user group having the highest perceived value to the provider are allocated keys from a first domain, and wherein keys from the first domain share fewer ancestors with keys from other domains than said keys from other domains share with each other.
4. A method according to claim 3 wherein keys from the first domain share only one ancestor with said keys from other domains.
5. A method according to claim 1 wherein the ancestral hierarchy has a binary tree architecture.
6. A method according to claim 1 wherein the policy provides that users are grouped according to a perceived susceptibility of them ceasing to require the service, and a first user group having the highest perceived susceptibility are allocated keys from a first domain, and wherein keys from the first domain share fewer ancestors with keys from other domains than said keys from other domains share with each other.

7. A method according to claim 6 wherein keys from the first domain share only one ancestor with said keys from other domains.
8. A method according to claim 1 wherein varying levels of service are available and a group of users of a low-service level are allocated placebo keys providing no security, thereby to obviate a need to reconfigure other user's keys upon their invalidation.
9. A method according to claim 8 wherein the service is a dynamic service and its value is ephemeral and based upon its contemporaneous nature.
10. A method of managing security key distribution to a plurality of users of a service comprising the steps of:
defining levels of service provision;
allocating keys to users which are indicative to a service provider of the level of service to which they are entitled; and
for at least one level of service provision, allocating placebo keys which do not provide security for the provision of the service.
11. A method according to claim 10 wherein the placebo keys operate in such a manner that a user is not able to perceive a difference between a functioning security key and a placebo key.
12. A method according to claim 10 wherein the service is dynamic and its value is ephemeral and based upon its contemporaneous nature.

MANAGEMENT OF SECURITY KEY DISTRIBUTION

ABSTRACT

5

Security keys for the provision of a secure service such as content provision are generated in an ancestral hierarchy, so that invalidation of a key in the hierarchy results in a need to reconfigure all other keys in the hierarchy to the extent they share common ancestry. When a user subscription to the service lapses, a decision on
10 invalidation of their key is based in a determination of whether it's more costly to the subscriber to invalidate the key, or continue providing an unpaid-for service. Keys can be allocated to users from domains of the hierarchy on the basis of their economic value to the provider, with higher value users being allocated keys from domains which share fewer common ancestors with other users of other domains than those
15 users share with each other, to minimise inconvenience to high value users of key reconfiguration.

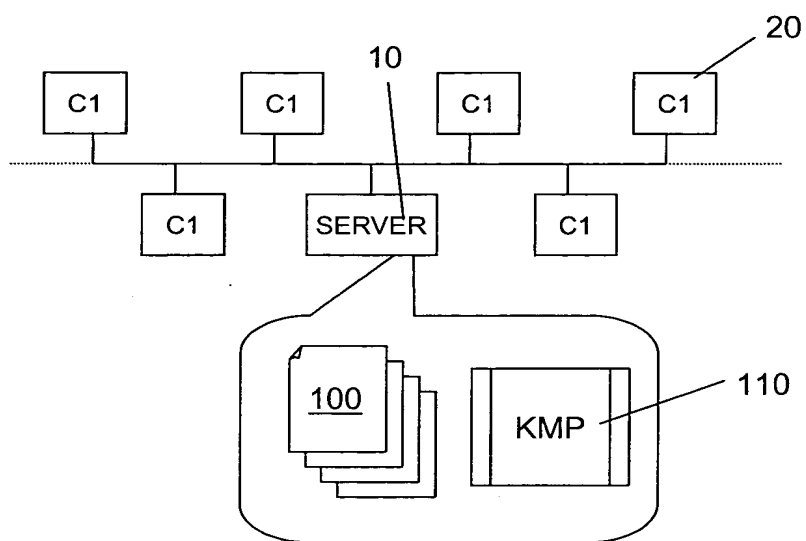


Fig. 1

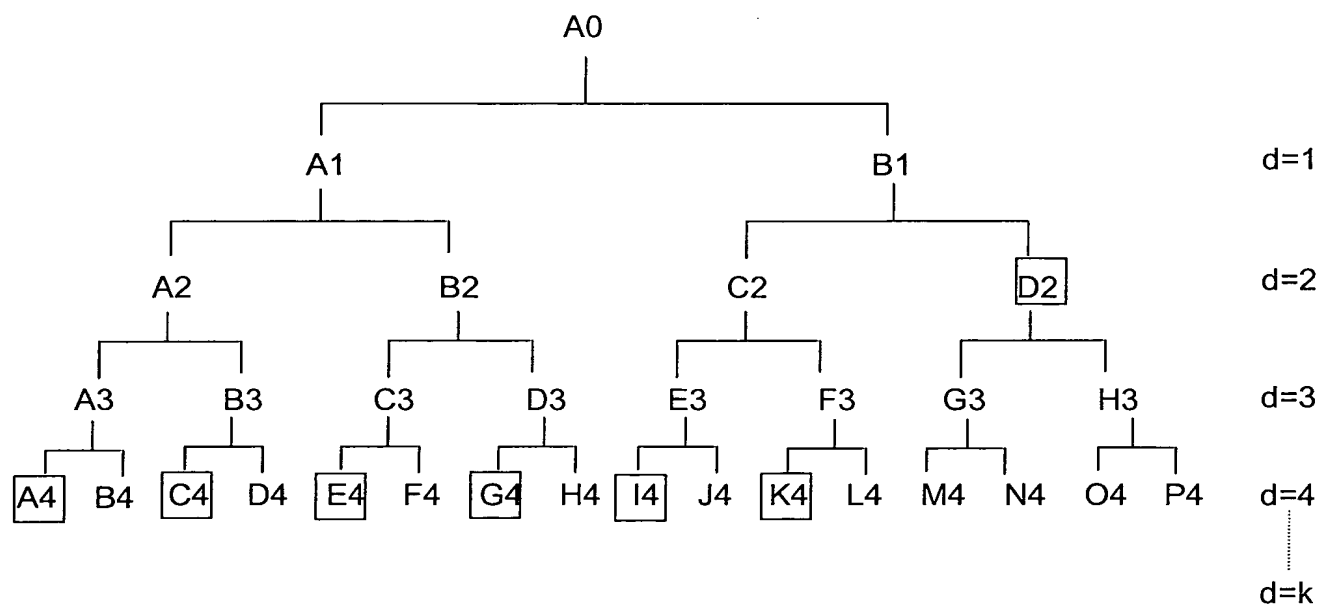


Fig. 2

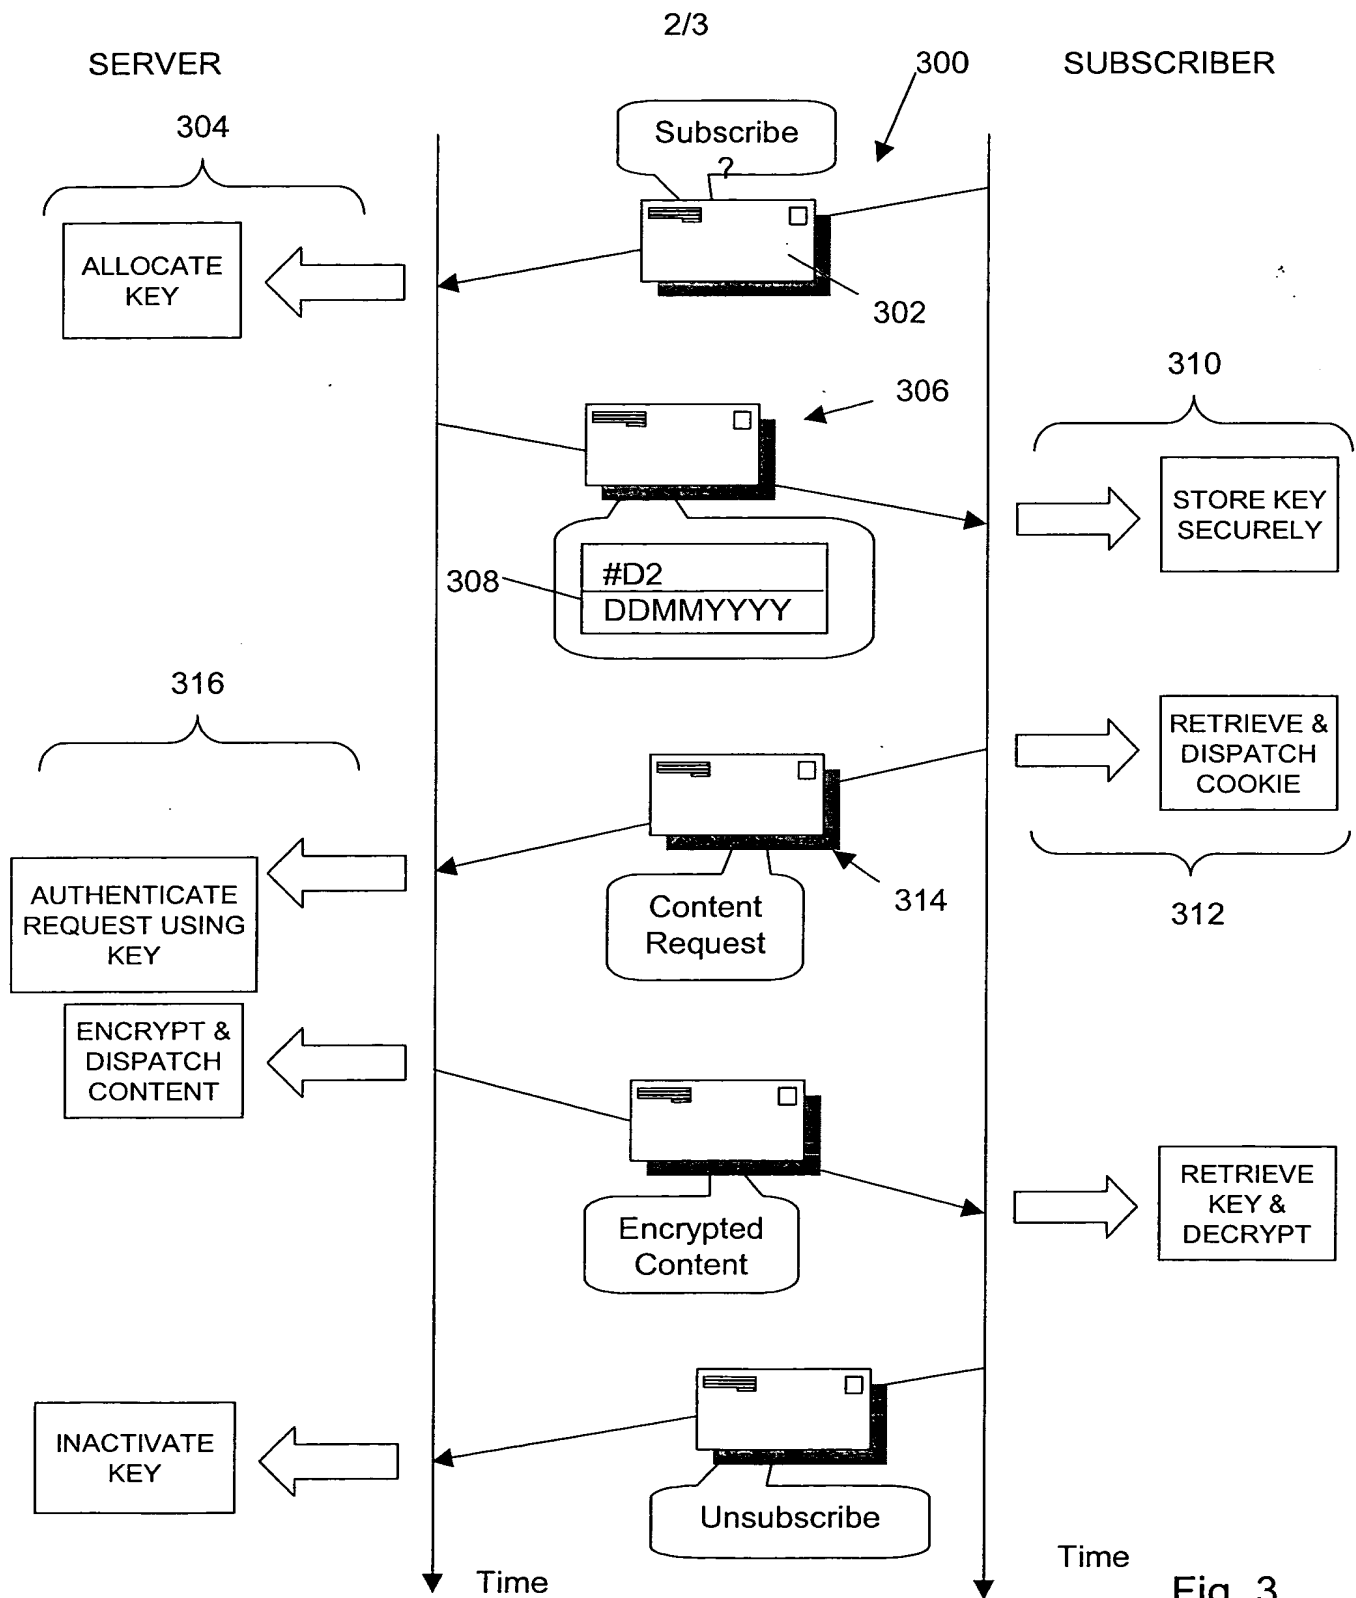


Fig. 3

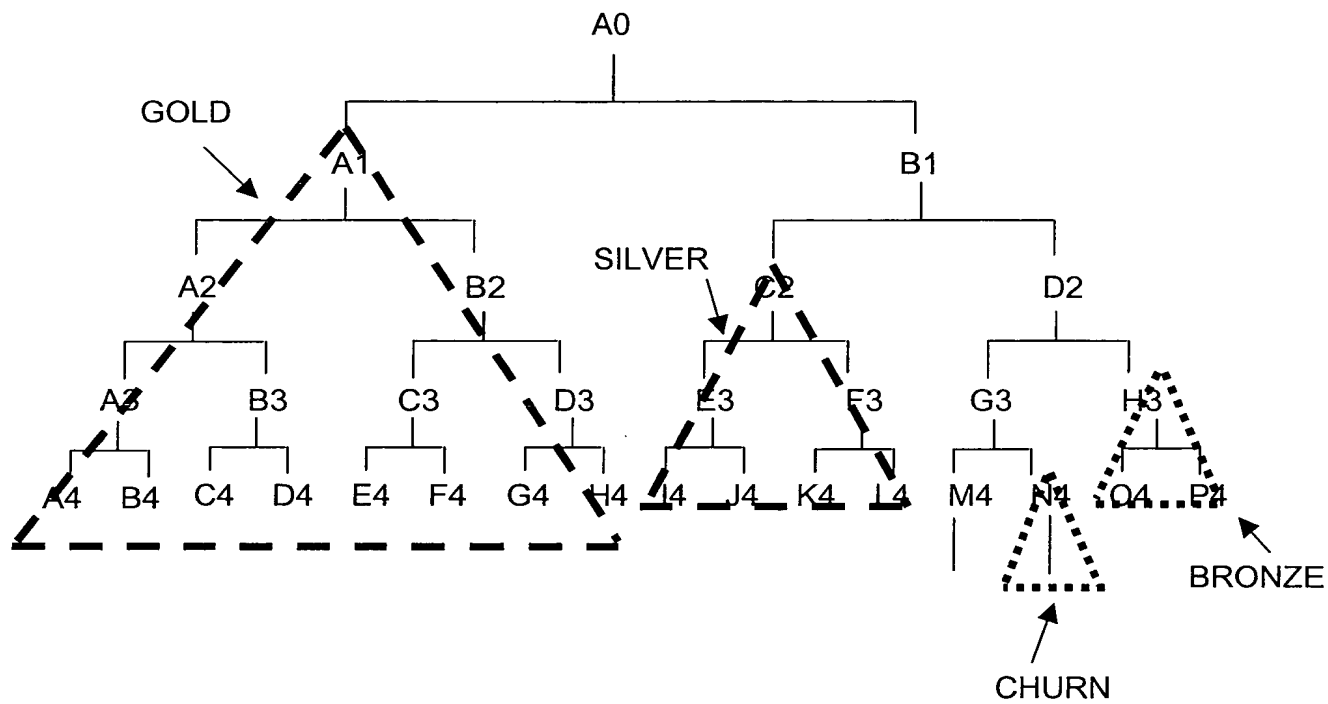


Fig. 4

HEWLETT-PACKARD
COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Antonio LAIN et al
084061-0544-200205652-2*